REMARKS/ARGUMENTS

This paper is being provided in response to the November 5, 2004 Office Action for the above-referenced application. In this response, Applicant has added new Claims 42-44, and amended Claims 1, 18, 27, 28, and 31-33, 35, 37, and 39-41 in order to clarify that which Applicant deems to be the claimed invention. Applicant respectfully submits that the amendments to the claims and the newly added claims are all supported by the originally filed application.

Applicant thanks the Examiner for the allowance of Claims 8-11.

The rejection of Claims 31-33 and 35-41 under 35 U.S.C. 101 has been addressed by amendments made herein. Accordingly, Applicant respectfully requests that the rejection be reconsidered and withdrawn.

The rejection of Claims 1-3, 5-19, 21-33, and 35-41 under 35 U.S.C. § 112, ¶1 has been addressed by amendments made herein in accordance with remarks set forth in the Office Action. Applicant has amended independent Claims 1, 18 and 31 to recite the feature of ***wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories***. Applicant respectfully submits that support for the foregoing feature may be found in the originally filed application. For example, the specification describes, at page 11, lines 7- 16, that in the case of data that has not been modified, the corresponding control data element indicates that the data has not been modified while, in the case of data that has been modified, the corresponding control data element

indicates that the data is write pending. The specification further states that the control data for the slots is written to both cache memories, 22, 24. The control data in one of the cache memories is identical to the control information in the other one of the cache memories. Thus, the specification states that control data is stored in both caches when the data has not been modified, as recited in Applicant's amended independent Claims 1, 18, and 31.

Accordingly, Applicant respectfully requests that the rejection be reconsidered and withdrawn.

The rejection of Claims 31-33, and 35-41 under 35 U.S.C. § 112, ¶1 has been addressed by amendments made herein in accordance with remarks set forth in the Office Action. In particular, Applicant has amended the claims to recite "processor executable" rather than "machine executable" features as suggested in the Office Action.

Accordingly, Applicant respectfully requests that the rejection be reconsidered and withdrawn.

The rejection of Claims 1-3, 5-19, 21-33, and 35-41 under 35 U.S.C. § 112, second paragraph has been addressed by amendments made herein in accordance with remarks set forth in the Office Action.

Accordingly, Applicant respectfully requests that the rejection be reconsidered and withdrawn.

The rejection of Claims 1-2, 5-7, 12-19, 21-32 and 35-37 under 35 U.S.C. § 103(a) as being unpatentable over Dewey et al. (U.S. Patent No. 5,724,501, hereinafter referred to as "Dewey") in view of Kurokawa et al. (U.S. Patent No. 6,571,350, hereinafter referred to as "Kurokawa") is hereby traversed and reconsideration thereof is respectfully requested. Applicant respectfully submits that Claims 1-2, 5-7, 12-19, 21-32 and 35-37, as amended herein, are patentable over the cited references.

Applicant's Claim 1, as amended herein, recites a method of managing data in a cache, comprising: providing a first cache memory containing data; providing a second cache memory containing data, wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said corresponding disk data has not been modified , said control data is replicated in said first and said second cache memories and in response to a request for data that is stored in both the first cache memory and the second cache memory, selecting which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique. Claims 2, 5-7, and 12-17 depend from Claim 1.

Applicant's Claim 18, as amended herein, recites a system for managing data in a cache comprising: a first cache memory containing data; a second cache memory containing data wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said

corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories independent of whether said corresponding disk data is included in both said first and said second cache memories; and cache selection hardware for selecting, in response to a request for requested data that is stored in both the first cache memory and the second cache memory, which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique. Claims 19 and 21-30 depend from Claim 18.

Applicant's Claim 31, as amended herein, recites a computer program product stored on a computer readable medium for managing data in a cache, comprising: processor executable code for providing a first cache memory containing data; processor executable code for providing a second cache memory containing data, wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories independent of whether said corresponding disk data is included in both said first and said second cache memories; and processor executable code for, in response to a request for data that is stored in both the first cache memory and the second cache memory, selecting which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique. Claims 32, and 35-37 depend from Claim 31.

Dewey relates to improvements in fault tolerant data processing systems and methods. (Col. 1, Lines 6-7). To facilitate the quick recovery of data lost as a result of a controller or cache failure, a method and apparatus for recovering mirrored data in the cache is presented. The method includes placing the memory module associated with a failed controller in a failed-over mode in which data is recovered in two stages. A portion of the memory contains a summary of all mirrored data, or metadata. The metadata provides a summary record of all data that was written to the cache from the host, but does not contain a record of any non-mirrored data in the cache, such as data that was written to the cache from the disk. (Col. 4, Lines 15-28). Read requests are handled differently than write requests. The write data is written in a mirrored fashion both to the data memory 20A and to the mirror memory 21A. (Col. 5, Lines 39-52). Figure 2 includes non-mirrored portions 50 and 52 which are reserved for data written from disk to the cache in response to a host read request. The mirrored memory is reserved for data written from host to cache for eventual storage on disk. Portions 58 and 60 store a metadata list of all data contained in the mirrored memory sections. The list does not include any information on the non-mirrored sections. (Col. 5, Line 61- Col. 6, Line 14). To facilitate quick recovery of data lost as a result of a controller of cache failure, the memory module associated with a failed controller is placed in a failover mode in which data is recovered in two stages using a battery backup. (Col. 4, Lines 15-21). Upon a controller failure, the metadata is first copied to a backup controller over a serial link. During a secondary recovery stage, the backup controller processes new host commands in the foreground and fetches mirrored data from the failed cache in the background. (Col. 4, Lines 29-37).

Kurokawa discloses a data handling system having a redundant storage configuration. (Col. 1, Lines 8-9). Kurokawa's Figure 1 includes duplicate data storage in separate storage

units SU(0) 16 and SU(1) 26. One of the storage units SU(0) and SU(1) performs as a master and the other as a sub storage in accordance with each address. (Col. 4, Lines 57-63; Figure 1). The master storage regions and sub storage regions are interleaved. (Col. 5, Line 65-Col. 6, Line 5; Figure 2). Kurokawa's Figure 8 includes a work storage WS unit in each of SU(0) and SU(1). Each WS is a cache memory smaller than the main storage which retains copies of a part of data in the main storages. Given a store or fetch request, the storage unit SU(0) or SU(1) stores or fetches data directly to or from the WS(0) or WS(1) when the WS(0) or WS(1) contains the desired data to be referenced. When the desired data is not found in WS(0) or WS(1), the storage SU(0) or SU(1) fetches a block of data including the desired data from the MS(0) or MS(1) and transfers the block in the WS. The storage unit SU(0) or SU(1) again accesses the WS(0) or WS(1) for a fetch or a store operation. (Col. 9, Line 60-Col. 10, Line 27; Figure 8).

Applicant's Claim 1, as amended herein, is neither disclosed nor suggested by the references in that the references, taken separately or in combination, neither disclose nor suggest ***a method of managing data in a cache, comprising: providing a first cache memory containing data; providing a second cache memory containing data, wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories; and in response to a request for data that is stored in both the first cache memory and the second cache memory, selecting which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique,*** as set forth in amended Claim 1. Dewey discloses handling write data in a mirrored

fashion writing the write data to two memories, and using a non-mirrored caching technique for data written to cache in response to a host read request. Dewey discloses a metadata list of data in mirrored memory sections. The metadata list does not contain any information on non-mirrored memory, such as data that is read by a host. In connection with a read operation, data is not modified and may be written from disk to cache in the non-mirrored memory. Thus, Dewey appears silent regarding any disclosure or suggestion of replicating control data when the corresponding disk data has not been modified. Kurokawa also appears silent with regard to any disclosure or suggestion of replicating control data when the corresponding disk data has not been modified. Accordingly, the references neither teach, disclose nor suggest at least the feature of ***wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories,*** as set forth in amended Claim 1.

The Office Action appears to render Dewey's metadata analogous to Applicant's recited "control data". Dewey's metadata, which is in both memories 20A and 21A, only includes information on mirrored memory for data written or modified by a host. Dewey's metadata does not include any information about non-mirrored memory for data of a host read request (e.g., data not modified). Dewey does not appear to make any suggestion or disclosure of replicating any data associated with non-modified cached data. Thus, Dewey cannot possibly teach anything about replicating control data for corresponding data in the cache that has not been modified.

For reasons similar to those set forth regarding Claim 1, Applicant's amended Claim 18 is also neither disclosed nor suggested by the references, taken separately or in combination in that the references neither disclose nor suggest ***a system for managing data in a cache comprising: a first cache memory containing data; a second cache memory containing data wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories independent of whether said corresponding disk data is included in both said first and said second cache memories; and cache selection hardware for selecting, in response to a request for requested data that is stored in both the first cache memory and the second cache memory, which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique,*** as set forth in amended Claim 18.

For reasons similar to those set forth regarding Claim 1, Applicant's amended Claim 31 is also neither disclosed nor suggested by the references, taken separately or in combination in that the references neither disclose nor suggest ***a computer program product stored on a computer readable medium for managing data in a cache, comprising: processor executable code for providing a first cache memory containing data; processor executable code for providing a second cache memory containing data, wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said corresponding disk***

data has not been modified, said control data is replicated in said first and said second cache memories independent of whether said corresponding disk data is included in both said first and said second cache memories; and processor executable code for, in response to a request for data that is stored in both the first cache memory and the second cache memory, selecting which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique, as set forth in amended Claim 31.

In view of the foregoing, Applicant respectfully requests that the rejection be reconsidered and withdrawn.

The rejection of Claims 3 and 33 under 35 U.S.C. § 103(a) as being unpatentable over Dewey in view of Kurokawa and further in view of Mason et al. (PCT/US98/19725, hereinafter referred to as "Mason") is hereby traversed and reconsideration thereof is respectfully requested. Applicant respectfully submits that Claims 3 and 33, as amended herein, are patentable over the cited references.

Claim 3 depends from Claim 1, and Claim 33 depends from Claim 31. For reasons set forth above, Claims 1 and 31 are neither disclosed nor suggested by Dewey and Kurokawa. For reasons set forth below, Applicant respectfully submits that combining Dewey and Kurokawa with Mason also neither discloses nor suggests Claims 1 and 31, and claims that depend therefrom.

Mason relates to mass storage systems in which stored logical volumes are duplicated in mirror form. (See page 1, lines 4-7). Mason describes dynamically adjusting the mirror policy for a disk drive system by periodically collecting statistics for the reading and writing of data to mirrored logical volumes and determining from the collected statistics whether the mirror service policy should continue or change. (See Abstract).

Applicant's Claim 1, as amended herein, is neither disclosed nor suggested by the references, taken separately or in combination, in that the references neither disclose nor suggest ***a method of managing data in a cache, comprising: providing a first cache memory containing data; providing a second cache memory containing data, wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories; and in response to a request for data that is stored in both the first cache memory and the second cache memory, selecting which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique,*** as set forth in amended Claim 1. For reasons set forth above, Dewey and Kurokawa do not disclose or suggest at least the feature of ***wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories,*** as set forth in amended Claim 1. Mason also appears silent with regard to the feature of ***wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories,*** as set forth in amended Claim 1.

Thus, Mason does not overcome the deficiencies of Dewey and Kurokawa with respect to Applicant's amended Claim 1. Accordingly, the references neither teach, disclose or suggest at least the feature of ***wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories,*** as set forth in amended Claim 1.

For reasons similar to those set forth regarding Claim 1, Applicant's Claim 31 is also neither disclosed nor suggested by the references, taken separately or in combination in that the references neither disclose nor suggest ***a computer program product stored on a computer readable medium for managing data in a cache, comprising: processor executable code for providing a first cache memory containing data; processor executable code for providing a second cache memory containing data, wherein at least some of the data contained in the first cache memory is the same as at least some of the data contained in the second cache memory, wherein data contained in said first and said second cache memories includes control data and corresponding disk data, wherein, when said corresponding disk data has not been modified, said control data is replicated in said first and said second cache memories independent of whether said corresponding disk data is included in both said first and said second cache memories; and processor executable code for, in response to a request for data that is stored in both the first cache memory and the second cache memory, selecting which one of the first and second cache memories to use to obtain the requested data in accordance with an access balancing technique,*** as set forth in amended Claim 31.

In view of the foregoing, Applicant respectfully requests that the rejection be reconsidered and withdrawn.

Applicant respectfully submits that newly added Claims 42-44 are also patentable over the cited references.

Based on the above, Applicant respectfully requests that the Examiner reconsider and withdraw all outstanding rejections and objections. Favorable consideration and allowance are earnestly solicited. Should there be any questions after reviewing this paper, the Examiner is invited to contact the undersigned at 617-248-4042.

Respectfully submitted,
CHOATE, HALL & STEWART

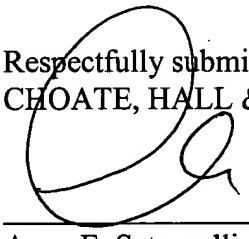

Anne E. Saturnelli
Registration No. 41,290

Patent Group
CHOATE, HALL & STEWART
Exchange Place
53 State Street
Boston, MA 02109-2804
Tel: (617) 248-5000
Fax: (617) 248-4000

Date: December 15, 2004